

EXHIBIT Y

Exhibit 17 – Lee

'156 Patent

Claim Limitation (Claim 7)	Exemplary Disclosure
[156a] A device comprising:	<p>Lee's, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i>, discloses a device. Specifically, Lee describes an FPGA-based face detector using neural networks. <i>See, e.g.</i>:</p> <p>“The study implemented an FPGA-based face detector using Neural Networks and a scalable Floating Point arithmetic Unit (FPU).” Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 315 (Abstract).</p> <p>“In Section 2, we describe our FPGA implementation of an FPGA-based face detector using neural networks.” Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 315.</p>
[156b] at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	<p>Lee discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See, e.g.</i>:</p> <p>“The FPU provides dynamic range and reduces the bit of the arithmetic unit more than fixed point method does. These features led to reduction in the memory so that it is efficient for neural networks system with large size data bits. The arithmetic unit occupies 39~45% of the total neural networks system area. Therefore bits reduction is needed not only for memory but also for a FPU and system size. Reduction from FPU 32 bits (IEEE 754 single precision) to 16 bits reduced the size of memory and arithmetic units by 50%, having only 1.25% deterioration in the detection rate. In order to determine the least and acceptable bits of the FPU, we examined how representation errors affect a detection rate through the MRRE. The scalable FPU and the error analysis may be useful to determine the details, especially area and speed of FPU for the embedded neural network system.” Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 315 (Abstract).</p> <p>“After face data come in the input node, they are calculated by MAC with weights. Face or non-face is determined by comparing output results and thresholds. For example, if output is</p>

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure
<p>[156c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and</p>	<p>larger than threshold, it means Face. Here, on the FPGA, this determination is decided easily by checking a sign bit after subtraction between output results and threshold.” Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 316.</p> <p>Lee discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See, e.g.:</i></p> <p>Lee discloses the use of a floating-point format due to its higher dynamic range:</p> <p>“Floating point numbers and fixed point numbers are representative number systems to express real numbers. Floating Point Unit (FPU) is slower and bigger than FiXed point Unit (FXU). However FPU provides a good dynamic range which is helpful for neural networks using wide number range with smaller bits. It is also useful to reduce memory having weights [4, 5].” Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 315.</p> <p>The floating point unit in Lee operates with reduced precision:</p> <p>2.3 Implementation of Reduced Precision FPU</p> <p>Fig. 4 shows the diagram of top module. The module consists of control logic and an FPU arithmetic unit.</p>

Exhibit 17 – Lee

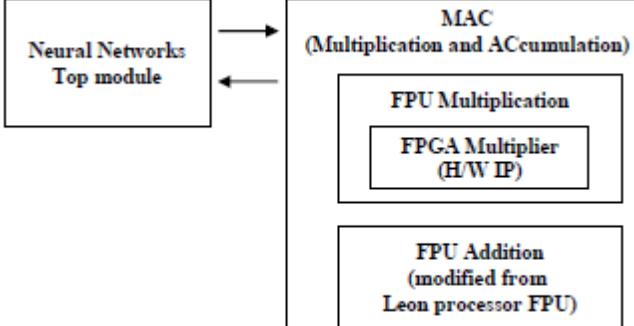
Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p data-bbox="925 612 1368 644">Fig. 4 Block diagram of FPU NN</p> <p data-bbox="692 693 1883 758">Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 316-17.</p> <p data-bbox="692 807 1698 840">The reduced precision level of Lee is dependent on the acceptable error range:</p>

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>3 Affection of Bit Reduction on FPU</p> <p>3.1 Representation Errors by FPU Bit Reduction</p> <p>The number of bits in FPU is important to the area and operating speed [8]. Therefore we need to decide the least number of bits within the acceptable error range. We used MRRE [9] as one of the indices of floating point arithmetic accuracy as shown in Table 1. MRRE is the Maximum Relative Representation Error, which is the relative distance between a real number and a represented number. The MRRE can be obtained as follows:</p> $MRRE = \frac{1}{2} \times ulp \times \beta \quad (5)$ <p>Where <i>ulp</i> is a unit in the last position and β is the exponent base.</p> <p>Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 317.</p>

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure																				
	<p>Specifically, Lee contemplates a 16-bit floating point unit with six exponent and bit nine mantissa bits:</p> <p style="text-align: center;">Table 1 MRRE of four Floating Point Units</p> <table border="1" data-bbox="783 385 1431 649"> <thead> <tr> <th data-bbox="783 385 903 425">Unit</th><th data-bbox="903 385 1015 425">β, e, m</th><th data-bbox="1015 385 1142 425">Range</th><th data-bbox="1142 385 1431 425">MRRE</th></tr> </thead> <tbody> <tr> <td data-bbox="783 425 903 483">FPU32</td><td data-bbox="903 425 1015 483">2, 8, 23</td><td data-bbox="1015 425 1142 483">$2^{2^8-1}=2^{255}$</td><td data-bbox="1142 425 1431 483">$0.5 \square 2^{-23} \square 2=2^{-23}$</td></tr> <tr> <td data-bbox="783 483 903 540">FPU24</td><td data-bbox="903 483 1015 540">2, 6, 17</td><td data-bbox="1015 483 1142 540">$2^{2^6-1}=2^{63}$</td><td data-bbox="1142 483 1431 540">$0.5 \square 2^{-17} \square 2=2^{-17}$</td></tr> <tr> <td data-bbox="783 540 903 597">FPU20</td><td data-bbox="903 540 1015 597">2, 6, 13</td><td data-bbox="1015 540 1142 597">$2^{2^6-1}=2^{63}$</td><td data-bbox="1142 540 1431 597">$0.5 \square 2^{-13} \square 2=2^{-13}$</td></tr> <tr> <td data-bbox="783 597 903 649">FPU16</td><td data-bbox="903 597 1015 649">2, 6, 9</td><td data-bbox="1015 597 1142 649">$2^{2^6-1}=2^{63}$</td><td data-bbox="1142 597 1431 649">$0.5 \square 2^{-9} \square 2=2^{-9}$</td></tr> </tbody> </table> <p><i>See Lee and Ko, An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit at 317.</i></p> <p>“The FPU representation error increase with MAC calculation of neural networks. The output difference can be expressed as following equations with the notation depicted in Fig. 6.</p>	Unit	β, e, m	Range	MRRE	FPU32	2, 8, 23	$2^{2^8-1}=2^{255}$	$0.5 \square 2^{-23} \square 2=2^{-23}$	FPU24	2, 6, 17	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-17} \square 2=2^{-17}$	FPU20	2, 6, 13	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-13} \square 2=2^{-13}$	FPU16	2, 6, 9	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-9} \square 2=2^{-9}$
Unit	β, e, m	Range	MRRE																		
FPU32	2, 8, 23	$2^{2^8-1}=2^{255}$	$0.5 \square 2^{-23} \square 2=2^{-23}$																		
FPU24	2, 6, 17	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-17} \square 2=2^{-17}$																		
FPU20	2, 6, 13	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-13} \square 2=2^{-13}$																		
FPU16	2, 6, 9	$2^{2^6-1}=2^{63}$	$0.5 \square 2^{-9} \square 2=2^{-9}$																		

Exhibit 17 – Lee

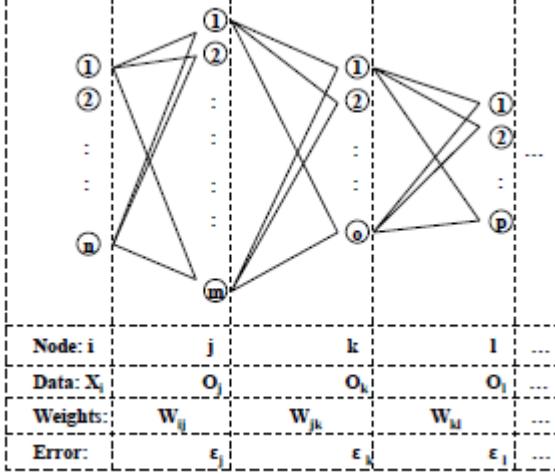
Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p data-bbox="925 739 1290 771">Fig. 6 Neural networks model</p> <p data-bbox="777 816 1410 848">The error of the 1st Layer can be described as (6).</p> $\begin{aligned} \varepsilon_j &= O^J_j - O_j \\ &= f(\sum_{i=1}^n W^f_{ij} X^f_{ij}) - f(\sum_{i=1}^n W_{ij} X_i) + \varepsilon_f. \end{aligned} \quad (6)$ <p data-bbox="777 1011 1417 1183">Where ε_j represents the hidden layer error, ε_k represents total error generated between hidden layer and output layer, and ε_f represents the non-linear function error. W represents the weights and O represents the output of hidden layer.</p> <p data-bbox="777 1188 1417 1258">The term W^f and X^f are new data including the finite precision error. Both are described by</p> $W^f = W + \varepsilon_W \text{ and } X^f = X + \varepsilon_X, \text{ respectively.}$

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure
	<p>Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 317.</p> <p>After putting the MRRE between FPU32 and other FPU bits into error terms in equations (16) and (18) using the MATLAB and real face data, we can finally find the total accumulated error of neural networks. We considered truncation for a rounding in this paper. Therefore a round-to-nearest scheme will reduce the error more [9].</p> <p>In order to find the detection rate errors, we modeled an EER (Equal Error Rate) graph using Gaussian distribution function and calculated max detection rate error by threshold error variation as shown in Fig. 7.</p> <p>Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 318.</p>

Exhibit 17 – Lee

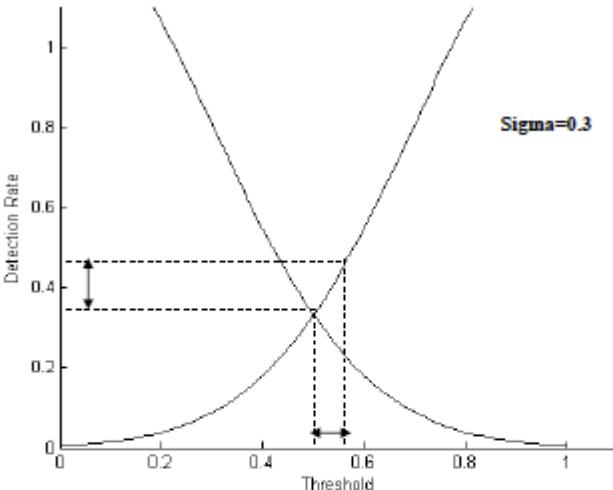
Claim Limitation (Claim 7)	Exemplary Disclosure
	 <p>The figure is a graph titled "Sigma=0.3" showing the relationship between Detection Rate (Y-axis) and Threshold (X-axis). The Y-axis ranges from 0 to 1.0 with increments of 0.2. The X-axis ranges from 0 to 1.0 with increments of 0.2. Two curves are plotted: a solid curve starting at (0,0) and increasing towards (1,1), and a dashed curve starting at (0,1) and decreasing towards (1,0). The two curves intersect at a threshold of approximately 0.5 and a detection rate of approximately 0.35. Dashed lines indicate these intersection points. A double-headed arrow on the Y-axis spans from 0.35 to 0.45, labeled "Detection Rate". A double-headed arrow on the X-axis spans from 0.1 to 0.5, labeled "Threshold".</p> <p data-bbox="908 750 1393 783">Fig. 7 EER graph of detection system</p> <p data-bbox="802 832 1499 946">Bit reduction of 50% from FPU32 to FPU16 is affected by only 1.25 % deterioration in the detection rate as shown in Table 2.</p> <p data-bbox="686 971 1879 1044">Lee and Ko, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 319.</p>

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure																				
	<p data-bbox="804 238 1199 270">Table 2 Detection rate errors</p> <table border="1" data-bbox="804 270 1501 556"> <thead> <tr> <th data-bbox="804 279 967 355">Unit</th><th data-bbox="967 279 1108 355">MRRE</th><th data-bbox="1108 279 1311 355">NN OUT Error (Max)</th><th data-bbox="1311 279 1501 355">Detection Rate Errors</th></tr> </thead> <tbody> <tr> <td data-bbox="804 355 967 393">FPU32</td><td data-bbox="967 355 1108 393">2^{-23}</td><td data-bbox="1108 355 1311 393">4.08E-7</td><td data-bbox="1311 355 1501 393">7.5E-7</td></tr> <tr> <td data-bbox="804 393 967 430">FPU24</td><td data-bbox="967 393 1108 430">2^{-17}</td><td data-bbox="1108 393 1311 430">2.61E-5</td><td data-bbox="1311 393 1501 430">4.8E-5</td></tr> <tr> <td data-bbox="804 430 967 468">FPU20</td><td data-bbox="967 430 1108 468">2^{-13}</td><td data-bbox="1108 430 1311 468">4.18E-4</td><td data-bbox="1311 430 1501 468">7.7E-4</td></tr> <tr> <td data-bbox="804 468 967 556">FPU16</td><td data-bbox="967 468 1108 556">2^{-9}</td><td data-bbox="1108 468 1311 556">0.0067</td><td data-bbox="1311 468 1501 556"><u>0.0125</u> <u>(1.25%)</u></td></tr> </tbody> </table> <p data-bbox="688 589 1892 659"><i>Lee and Ko, An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit at 319.</i></p> <p data-bbox="688 698 1913 809"><i>“It is confirmed through MATLAB simulation that this polynomial approximation made about 5 % error in a detection rate.” Lee and Ko, An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit at 316.</i></p> <p data-bbox="688 845 1913 988"><i>“We also showed that bits reduction from FPU 32 bits to 16 bits reduced the size of memory and arithmetic units by 50%, having only 1.25% deterioration in the detection rate.” Lee and Ko, An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit at 320.</i></p> <p data-bbox="688 1029 1860 1099"><i>See also Appendix to Responsive Contentions Regarding Non-Infringement and Invalidity (“Responsive Contentions”) (detailing error rates associated with different mantissa sizes).</i></p>	Unit	MRRE	NN OUT Error (Max)	Detection Rate Errors	FPU32	2^{-23}	4.08E-7	7.5E-7	FPU24	2^{-17}	2.61E-5	4.8E-5	FPU20	2^{-13}	4.18E-4	7.7E-4	FPU16	2^{-9}	0.0067	<u>0.0125</u> <u>(1.25%)</u>
Unit	MRRE	NN OUT Error (Max)	Detection Rate Errors																		
FPU32	2^{-23}	4.08E-7	7.5E-7																		
FPU24	2^{-17}	2.61E-5	4.8E-5																		
FPU20	2^{-13}	4.18E-4	7.7E-4																		
FPU16	2^{-9}	0.0067	<u>0.0125</u> <u>(1.25%)</u>																		

Exhibit 17 – Lee

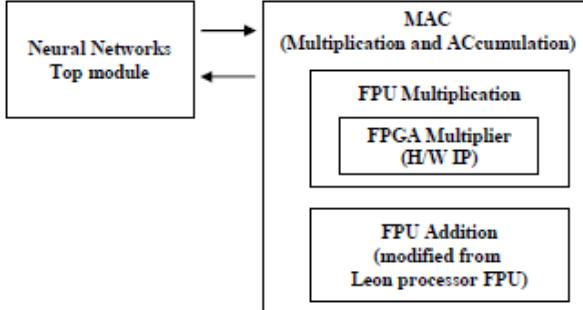
Claim Limitation (Claim 7)	Exemplary Disclosure
[156d] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit;	<p>Lee discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. <i>See, e.g.</i>:</p> <p>2.3 Implementation of Reduced Precision FPU</p>
	<p>Fig. 4 shows the diagram of top module. The module consists of control logic and an FPU arithmetic unit.</p>  <p>Fig. 4 Block diagram of FPU NN</p> <p>Lee, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 316-17.</p>
[156e] wherein the at least one first computing device comprises at least one of a central processing unit (CPU), a graphics processing unit (GPU), a field programmable gate array (FPGA), a microcode-based processor, a hardware sequencer, and a state machine;	<p>To the extent that Singular contends that Lee does not itself identify one of the disclosed computing devices, notwithstanding this disclosure, use of said computing device would have been obvious based on Lee alone or in combination with other disclosed prior art, including without limitation Lee, Belanovic, GRAPE-3, and Cray T3d, for the reasons explained in the Responsive Contentions.</p>

Exhibit 17 – Lee

Claim Limitation (Claim 7)	Exemplary Disclosure																				
<p>[156f] and, wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.</p>	<p>“The arithmetic unit occupies 39~45% of the total neural networks system area. Therefore bits reduction is needed not only for memory but also for a FPU and system size. Reduction from FPU 32 bits (IEEE 754 single precision) to 16 bits reduced the size of memory and arithmetic units by 50%, having only 1.25% deterioration in the detection rate.” Lee, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 315 (abstract).</p> <p>“Bit reduction of the FPU led to an area reduction and the faster operating clock speed. 50% of bit reduction from FPU 32 to FPU 16 resulted in 50% of addition area (250/486) and memory reduction (1880/3760) as shown in Table 4.” Lee, <i>An FPGA-Based Face Detector Using Neural Network and a Scalable Floating Point Unit</i> at 319.</p> <p>Table 4 Analysis of 32, 24, 20, 16-bit of FPUs</p> <table border="1" data-bbox="709 711 1326 1070"> <thead> <tr> <th data-bbox="709 711 946 817">Arithmetic Unit</th><th data-bbox="946 711 1094 817">Memory (Kbits)</th><th data-bbox="1094 711 1199 817">NN Area (Slices)</th><th data-bbox="1199 711 1326 817">FPU Adder Area (Slices)</th></tr> </thead> <tbody> <tr> <td data-bbox="709 817 946 866">FPU 32</td><td data-bbox="946 817 1094 866">3760</td><td data-bbox="1094 817 1199 866">1077</td><td data-bbox="1199 817 1326 866">486 (45%)</td></tr> <tr> <td data-bbox="709 866 946 940">FPU 24</td><td data-bbox="946 866 1094 940">2820 (75%)</td><td data-bbox="1094 866 1199 940">878</td><td data-bbox="1199 866 1326 940">403 (45%)</td></tr> <tr> <td data-bbox="709 940 946 1013">FPU 20</td><td data-bbox="946 940 1094 1013">2350 (63%)</td><td data-bbox="1094 940 1199 1013">750</td><td data-bbox="1199 940 1326 1013">300 (40%)</td></tr> <tr> <td data-bbox="709 1013 946 1070">FPU 16</td><td data-bbox="946 1013 1094 1070">1880 (50%)</td><td data-bbox="1094 1013 1199 1070">650</td><td data-bbox="1199 1013 1326 1070">250 (39%)</td></tr> </tbody> </table> <p>To the extent that Singular contends that Lee does not identify a device with at least 100 multiplication execution units, notwithstanding this disclosure, such a device would have been obvious given intervening FPGA developments for the reasons explained in the Responsive Contentions.</p>	Arithmetic Unit	Memory (Kbits)	NN Area (Slices)	FPU Adder Area (Slices)	FPU 32	3760	1077	486 (45%)	FPU 24	2820 (75%)	878	403 (45%)	FPU 20	2350 (63%)	750	300 (40%)	FPU 16	1880 (50%)	650	250 (39%)
Arithmetic Unit	Memory (Kbits)	NN Area (Slices)	FPU Adder Area (Slices)																		
FPU 32	3760	1077	486 (45%)																		
FPU 24	2820 (75%)	878	403 (45%)																		
FPU 20	2350 (63%)	750	300 (40%)																		
FPU 16	1880 (50%)	650	250 (39%)																		

Exhibit 17 – Lee

'273 Patent

Claim Limitation (Claim 53)	Exemplary Disclosure
[273a] A device:	Lee discloses a device. <i>See [156a].</i>
[273b] comprising at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Lee discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b].</i>
[273c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X % of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input;	Lee discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=5% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c].</i>
[273d] wherein the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide.	Lee discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. <i>See [156f]</i>

Exhibit 17 – Lee

'961 Patent

Claim Limitation (Claim 4)	Exemplary Disclosure
[961a] A device comprising:	Lee discloses a device. <i>See [156a].</i>
[961b] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Lee discloses at least one first low precision high dynamic range (LPHDR) execution unit adapted to execute a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [156b].</i>
[961c] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input; and	Lee discloses the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input. <i>See [156c].</i>
[961d] at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit.	Lee discloses at least one first computing device adapted to control the operation of the at least one first LPHDR execution unit. <i>See [156d].</i> .

Exhibit 17 – Lee

Claim Limitation (Claim 4)	Exemplary Disclosure
Claim Limitation (Claim 13)	Exemplary Disclosure
[961e] A device comprising:	Lee discloses a device. <i>See [961a].</i>
[961f] a plurality of components comprising:	Lee discloses a plurality of components. <i>See [961b] + [961d].</i>
[961g] at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value,	Lee discloses at least one first low precision high-dynamic range (LPHDR) execution unit adapted to execute a first operation on a first input signal representing a first numerical value to produce a first output signal representing a second numerical value. <i>See [961b].</i>
[961h] wherein the dynamic range of the possible valid inputs to the first operation is at least as wide as from 1/1,000,000 through 1,000,000 and for at least X=10% of the possible valid inputs to the first operation, the statistical mean, over repeated execution of the first operation on each specific input from the at least X% of the possible valid inputs to the first operation, of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input differs by at least Y=0.2% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input.	Lee discloses the number of LPHDR execution units in the device exceeds by at least one hundred the non-negative integer number of execution units in the device adapted to execute at least the operation of multiplication on floating point numbers that are at least 32 bits wide. <i>See [961c].</i>